

"METHOD AND DEVICE FOR THE WFQ STATISTICAL MULTIPLEXING OF ATM
FLOWS"

5 Field of the invention

The present invention relates in general to communication systems employing the digital technique of cell quick switching, called ATM (Asynchronous Transfer Mode), used for transport of voice, video signals and data.

More in particular, the invention relates to a method and a device for the 10 statistical multiplexing in Weighted Fair Queuing (WFQ) technique of flows conveying ATM traffic.

The ATM technique has an ever increasing important role in the integrated 15 switching of flows of digital signals belonging to services for the transmission of voice, video and data signals, with different bandwidth requirements and differentiated traffic characteristics.

The ATM technique foresees in particular, that the information relevant to the 20 different services is organized in adjoining units having fixed length of 424 bits (53 bytes), called cells. In addition to the true data, said cells contain a header, carrying, among other things, the information necessary to the routing of the cell itself through the geographical network.

The connection network, which has the task to spatially switch the cells from an 25 input port to an output port, must be capable of developing high traffic volumes, in the range of some hundreds of Gbit/s.

The introduction of service priority mechanisms on the flow of ATM cells, be 30 they represented by single connections or groups of connections, is connected to the support in the network of different classes of quality of service (QoS).

The necessity to give guarantees for delay-sensitive sources and, at the same time support to data applications that, while tolerating longer delays, can employ band 35 statistical allocation principles, leads to the need to define and implement opportune integration schemes in the network. It must be added the possibility to introduce service classes at available band, or ABR, (Available Bit Rate) enabling to better employ the network resources and thus offer a more economic service to the users.

Background art

The article by Uwe Briem, Eugen Wallmeier, Christoph Beck, Fred Matthiesen, 35 under the title "TRAFFIC CONTROL FOR ATM SWITCH WITH PER VC QUEUING:

CONCEPT AND IMPLEMENTATION", published following the Proceedings of XVI World Telecom Congress, year 1997, describes a particular architecture of statistical multiplexers combined to an ATM switch, capable of better satisfying the above mentioned requirements imposed by the Traffic Management. Fig.1 of the article, coinciding with fig.1 of the present application, shows the architecture commonly found in the nodes of the ATM network and gives a starting base for the development of concepts useful to the description of the present invention.

Making reference to fig.1, we notice some blocks LIC₁, ..., LIC_r (Line Interface Cards) that represent cards of a line interface, each one operating at its own bit-rate, organized in groups canalizing the relevant flows towards the inputs of a same number of ATM multiplexers, denoted AMX₁, ..., AMX_k. These last, represent high speed internal interfaces towards the relevant statistical multiplexers, SMU₁, ..., SMU_k, which receive input flows LINK-IN and supply output flows LINK-OUT directed towards the inputs of a switching matrix ATM, denoted ASN (ATM Switching Network). A completely specular annotation applies to the section placed on the exit side of said matrix. In this case the flows coming out from the ASN matrix represent a same number of input flows LINK-IN for identical statistical multiplexers, SMU₁, ..., SMU_k, and the flows LINK-OUT coming out from these last, reach the ATM multiplexers (demultiplexers) AMX₁, ..., AMX_k, each of them shunting the traffic conveyed towards a relevant group of Line Interface Cards, LIC₁, ... LIC_r operating at different bit-rates.

In the operation, it is assumed that the ASN matrix is self-addressing and not locking. As it is already known, the main task of the statistical multiplexers is to organize the traffic at output, on the basis of the traffic variability at input, dynamically sharing the band available on the output link among the different active sessions, to the explicit purpose to maximize the traffic that can be globally controlled by the ASN matrix, assuring at the same time, the minimum band agreed in the Traffic Management protocol negotiation phase, to the different sessions. These targets are reached through a scheduling function, while a shaping functionality limits at top the band of the different connections, limiting the loss of ATM cells due to congestion.

To cope with these needs, the SMU multiplexers are fit with large dimensioned buffers, and therefore capable of adequately treating the channels conveying the non real-time traffic (nrt), such as for instance that defined in service classes nrt-VBR (nrt Variable Bit Rate), ABR (Available Bit Rate), and UBR (Unspecified Bit Rate). On the other hand, the channels conveying the real-time traffic, such as for instance that defined in CBR (Constant Bit Rate) and rt-VBR service classes do not require, of

course, large buffers.

Figures 2 and 3 of the present application are high level schematisations of the statistical multiplexers SMU, valid in general. Fig. 2 gives a functional block diagram, indicating how the ATM cells forming an input link LINK-IN are shared into n transmission queues QU1, ..., QUn and stored in a buffer controlled by a Buffer Managing BM function. Each queue, gathers cells coming from different elementary flows, forming a same number of connections within the network, supported by relevant virtual channels VC (Virtual Channel). The latter are conveniently shared into different virtual transmission paths VP (Virtual Path). The grouping of VCs according to transmission queue can take place on the basis of different criteria, for instance, according to the bit-rate of the line interface LIC, or according to the VP path within the network.

Fundamental for the explanation of the above mentioned BM function is the fact to keep updated the knowledge of the memory locations of the n queues and of the cells present in each queue. To this purpose, the Buffer Manager BM, updates at each cell acquired, a table of indicators, Q-ID1, ..., Q-IDn of the address of the last cell occupied in each cell queue (assumed as known the piling method within the queue). To perform the above mentioned operation, the BM block analyses the identification of the connection present in the Header portion of the ATM cell, in particular in VC (Virtual Channel) and VP (Virtual Path) fields, and employs this information to calculate an indicator Q-ID at the transmission queue. The acquisition or the rejection of a cell is decided on the basis of the traffic class and of the degree of occupation of the queue and of the buffer as a whole. In case of acceptance, the BM requires the writing of the cell in the buffer and decides whether the indicator Q-ID has to be stored or not, for the additional functions ending with the service of the queue and the cell transmission.

Downstream the BM block there is a SORTER block performing a service function, called Cell Scheduling, consisting in conveniently selecting ATM cells to place in sequence on the output link LINK-OUT. Multiple are the algorithms developed in order to accomplish the above-mentioned function; among these, the Weighted Fair Queuing one has particular importance to the purposes of the invention, hereinafter called WFQ.

Fig. 3 highlights the circuit setting of multiplexers SMU of fig. 2; a description shall be given below of the BM and SORTER blocks, at a good circuit detail level.

Reverting to the above mentioned article, and to fig. 3 in particular, shown more in detail in fig. 4 of the present application, it can be noticed the technological evolution

which a single statistical multiplexer SMU is subjected to, versus the generic functional representation of fig. 2, irrespective of the fact that it is placed at the input or at the output of the ASN switching matrix of fig. 1. In the representation of fig. 4, in fact, it is evident the neat separation operated among the transmission queues supporting real-time traffic and the queues supporting non real-time traffic. It is also clear the presence of a modular structure in which logic queues of QU cells of non real-time type, and the queues of the relevant queue indicators Q-ID, are grouped in m logic blocks identified by relevant indicators B-ID₁, ..., B-ID_m, while QU cells of real-time type are grouped in a single queue (CBR) visible at the lower end of the figure.

Each B-ID block includes its own sub-block SB (Scheduler Block) performing the scheduling function on the n logic queues Q-ID₁, ..., Q-ID_n of the block itself, and gives a sequence of indicators to an additional scheduling block SBS (Scheduler Block Scheduler) common to all the B-ID blocks and to the queue of real-time connections. The SBS scheduler reassembles the two kinds of traffic in the output flow LINK-OUT.

In the operation, the queues of ATM cells of each B-ID block, share a configurable and deterministic cell-rate, in other words this means that each block B-ID is assigned a fix band, possibly varying from block to block. The criteria for the assignment of the logic queues to the single blocks are indicated in the mentioned article. To generate the services of B-ID blocks, the scheduler SBS employs a calendar in which the B-ID indicators are continuously inserted and extracted according to an appropriate known criterion. Likewise, the scheduler SB of a generic block B-ID generates the services of the single queues with the aid of a calendar containing the Q-ID indicators of queues belonging to the B-ID block considered. When a block is served, the relevant scheduler SB determines the queue having right to the extraction of a cell from the buffer. The method according to which calendars are used, depends on one of the two scheduling methods foreseen, and in particular the method called Rate Shaping, to limit the peak cell-rate, and the WFO method to assure the minimum band agreed in the Traffic Management protocol negotiation phase, to non real-time connections.

For what said above, we can conclude that the structure described above, operates the band subdivision into queues belonging to different logic blocks through a set of m+1 calendars physically separated and organized in a double hierarchical level. This architecture presents the two following limitations:

- Impossibility to dynamically share the aggregate band among the blocks, since each block has right to a fix band, therefore, any possible "transmission gaps" of a

block cannot be used by other blocks.

- Impossibility to discretionary increase the number of blocks, since the circuit realization of a great number of BS schedulers (one scheduler per block) with calendar mechanism, which in practical applications can easily assume large dimensions, finds a limit set by the technology of the integrated semi-custom circuits.
- Need to perform a scanning of the calendar of a BS scheduler at each cell time to determine la queue of cells to serve, this involves the execution of number of operations in a short time, requiring the use of very fast and therefore expensive circuits.

The last two limitations of the known art in fig.4 partly fall, thanks to an invention developed in the laboratories of the Applicant, and already forming the object of an Italian patent application filed under number MI97A 002147 on 23-9-1997. This invention, shown in figures 5 and 6, enables to dynamically share the band among the transmission queues served by a generic block of queues through an original application of the WFQ algorithm. Unfortunately, the impossibility to dynamically share the band among the blocks persists.

Fig. 5 shows the functional arrangement of the blocks of a statistical multiplexer realized according to the invention just mentioned. Also in this case the transmission queues of non real-time connections are shared in m functional blocks addressed by the B-ID₁, ..., B-ID_m block indicators; the Q-ID indicators at single queues are submitted to a similar sharing. The indicators Q-ID destined to a generic block are therefore processed by an optional block SHP/SKD performing the shaping and scheduling functions to limit the peak band; after that, at any new occurrence of an indicator, according to the typology of the cell queue which the same is associated with, it is inserted in one of the FIFO lists (First In First Out), HP-FIFO, LP-FIFO, and VLP-FIFO, respectively, indicated in service priority order. For the sake of brevity, we did not indicate of a flag (traffic light) placed at the service of a queue grouping all the connections served in real-time, therefore bypassing the WFQ function, and it must be pointed out that the VLP-FIFO lists are optional. The FIFOs are used to perform the WFQ function according to the method described in claim 1, that is:

"Method for the management of resources in ATM technique for Weighted Fair Queuing applications in a system for the transmission on an output link of flows of ATM cells coming from an input link belonging to service classes of different quality, to at least some of said flows being associated a relevant guaranteed minimum band

assigned in the phase of service establishment, and being assured to each one of said cell flows the possibility to avail of an additional band possibly available on said output link, characterized in that it assigns to each one of said flows, in a renewal period multiple integer of the transmission time of a cell on said output link, a relevant number 5 of access rights to said output link, or tokens, proportional to said minimum band, and, for each renewal period, to include the following steps:

- inserting the cell of each one of said flows in a relevant transmission queue;
- assigning to each one of said transmission queues its own token number;
- inserting indicators of said transmission queues in a high priority storing structure;
- 10 - scanning said high priority storing structure serving the transmission queues time by time indicated, decreasing the tokens at their disposal and transferring the indicators of the queues that have finished their tokens inside a low priority storing structure;
- once the scanning of said high priority storing structure is completed, scanning said low priority storing structure serving the queues indicated by the same up to the end 15 of the renewal period".

As it can be noticed, the claim tenor does not reveal at all the possibility to subdivide the cell queues into functional blocks. In the text, this possibility is given for granted, since it is only vaguely referred to. In any case, the description of the means sorting the service among the different B-ID blocks is missing, and the reader can reasonably assume that they are tacitly similar to those known up to that time, that is of 20 the type of the SBS block of fig. 4. The latter is represented in fig. 5 by a SHP/SKD block that receives to the purpose the B-ID indicators.

Fig. 6 highlights the strict system-related aspect of the statistical multiplexer represented in fig.5; this figure shows a reference point for the comparison with successive similar figures that shall depict the scope of the invention of the present application, and of some embodiments. As it can be noticed, the buffer manager BM supplies the B-ID block and Q-ID queue indicators to the SHP/SKD-Q block, and the sole B-ID block indicator to the SHP/SKD-B block. These two blocks, as already said, 25 perform the shaping and scheduling functions for the service supplied to the blocks (SHP/SKD-B) and to the queues inside the blocks (SHP/SKD-Q). The indicators Q-ID extracted from the SHP/SKD-Q block are time by time written in one of the three FIFOs (fig.5) of a corresponding group of FIFOs included in a WFQ-Q block, which carry out the above mentioned WFQ function availing of a B-ID indicator extracted from the SHP/SKD-B block to address the group of FIFOs served. The indicators Q-ID extracted 30 from the WFQ-Q block are finally supplied to the BM block to complete the service. The indicators Q-ID extracted from the WFQ-Q block are finally supplied to the BM block to complete the service.

function on the output link LINK-OUT.

The examination of figures 5 and 6, clearly highlights the residual limitation of the above mentioned invention, due the fact that a function enabling to dynamically vary the band assigned to the blocks identified by indicators B-ID, is missing.

- 5 Furthermore, as it occurs in the practice, whenever the optional blocks SHP/SKD are actually present, partly re-emerge also the limitations of the known art of fig. 4 partly reappear, otherwise overcome, since the same architecture at m+1 calendars is newly proposed.

Scope of the invention

10 Therefore, scope of the present invention is to overcome the drawbacks of the statistical multiplexers for ATM applications realized through a functional block structure of cell queues according to the known art. More in particular, firstly the invention aims at overcoming the limitation due to the incapability to dynamically share the band available among the blocks, in order to better satisfy the needs of the actual traffic, and secondly the typical restrictions of scheduler array structures, that is the 15 complexity of their implementation.

A further scope of the present invention is to indicate a method for the statistical multiplexing in ATM technique enabling to better make use of the available band without creating congestion and drastically reducing the total number of operations 20 necessary to attain these purposes.

Summary of the invention

To attain these targets, scope of the present invention is a statistical multiplexing method of flows of ATM cells originated by connections having different quality service classes, said flows generating transmission queues of cells shared in functional blocks, 25 and at least some of said flows being assigned a minimum guaranteed band on the multiplexed flow and the possibility to share a residual band possibly available on the same flow through a known functionality, hereinafter called Weighted Fair Queuing, or WFQ, which dynamically assigns opportune quantities of tokens, evaluated on statistical basis, to single transmission queues, that spend said tokens to have access 30 to the residual band on the multiplexed flow (LINK-OUT), characterized in that said WFQ function includes dynamic assignment steps of appropriate quantities of tokens, evaluated on statistical basis, also to said functional blocks, that spend said tokens to enable their transmission queues to have access to the sharing of the residual band on the multiplexed flow (LINK-OUT), thus sharing the residual band on the multiplexed 35 flow among said functional blocks, as described also in claim 1.

An additional scope of the invention is a statistical multiplexing method having all the characteristics of the method described in claim 1, and the additional feature to perform the following at each cell time:

- a first scheduling of emission intervals of the indicators of the functional blocks to serve, suitable to expand said intervals to limit the aggregate peak band of the group of cell queues belonging to the blocks served;
- a second scheduling of emission intervals of the indicators of transmission queues to serve, suitable to expand said intervals to limit the peak band of queues served, as described in an appended claim of the method.

Another scope of the invention is a statistical multiplexer implementing the method described in claim 1, as described in the independent claim 13.

Another scope of the invention is a statistical multiplexer having all the means of the multiplexer of claim 13, and additional scheduling means of the service intervals of the functional blocks and relevant transmission queues, as described in an appended claim for the device.

Additional scopes of the invention are some embodiments of method and device, as they result from the relevant appended claims.

Advantages of the invention

A first considerable advantage of the present invention is due to the fact that the application of the WFQ algorithm, both at block level and at queue level, enables to extend the dynamic sharing of the band also to the block level, permitting to avail at best of the transmission resources. This assumed, an additional resulting advantage can be obtained grouping all the connections of service real-time classes of a block, for instance CBR, placing them in a high priority queue inside the block itself. In this way, the WFQ algorithm applied to the queues of the block considered, assures the minimization of the variation of cell delay, or CDV (Cell Delay Variation) caused to the CBR connections by the other connections belonging to the block. However, the delay of CBR connections is affected by the delay in the service of the block they belong to. Whenever the selection of the system parameters makes a problem of this dependence, it is possible to group the CBR connections of all the blocks in one high priority single block dedicated to the same. As it can be noticed, the invention enables to create in an extremely natural manner two different bypass levels, particularly suitable for the processing of real-time transmission flows.

A second important advantage of the statistical multiplexer of the present invention is that it needs only two calendars to perform the peak band control, both at

block and at queue level, limiting the band to both levels of the logic hierarchy. A first calendar is for the blocks, and a second one is for the queues; this last is supplied by the block queues selected time by time. Therefore, compared to the architectures of the known art, which replicate the calendars as many times as the blocks are, the logic
5 necessary to implement the peak rate control suffers a drastic reduction in complexity. At equal number of ports of the statistical multiplexer, it results a reduction in the size of the integrated circuits implementing the calendars and a consequent reduction of production costs. Furthermore, the removal of the technological restriction on the maximum number of blocks imposed by the replica of calendars, enables the
10 manufacturing of systems with a number of ports considerably higher than that possible up to date, ranging around one hundred. The architecture of the statistical multiplexer according to the present invention enables to impose the peak limitation in a selective way, since queues can subsist to which the pure WFQ algorithm is applied, that is with peak rate not limited, and queues limited at top in the band. This property, which applies also to the blocks, can be obtained through the configuration of appropriate
15 parameters via software.

Brief description of the drawings

The invention, together with further objects and advantages thereof, may be understood making reference to the following detailed description, taken in conjunction
20 with the accompanying drawings, in which:

- fig. 1 shows an architecture commonly found in ATM network nodes;
- figures 2 and 3 shows high level diagrams of the statistical multiplexers SMU of fig.
25 1 generally valid;
- fig. 4 shows a functional diagram of the SMU block of fig. 2, referred to the known art;
- fig. 5 shows an additional functional diagram of the SMU block of fig. 2, as it results from a previous patent application filed by the same Applicant;
- fig. 6 shows a functional block diagram of the SMU block of fig. 3 when the
functional diagram of fig. 5 applies;
- fig. 7 shows a functional diagram of the SMU block of fig. 2 realized according to the
present invention;
- fig. 8 shows a functional diagram of a SHP/SKD block of fig. 7;
- fig. 9 shows a functional block diagram of the SMU block of fig. 3 when the
functional representation of fig. 7 according to the invention, applies;
- 35 - fig. 10 shows a circuit diagram of a BM block of fig. 9;

- fig. 11 shows a circuit diagram valid for both the WFQ-Q and WFQ-B blocks of fig. 9;
- fig.12 shows in detail a generic TK-CNTj block and a second TK-CNT-B block of fig. 11;
- figures 13 through 16 show the flow charts representing the operation of the WFQ-B
5 block of fig.9;
- fig.17 shows a functional block diagram of an embodiment of the block diagram of fig.9;
- fig.18 shows a flow chart, representing the operation according to the embodiment
10 of fig.17; and
- figures 19 and 20 shows functional block diagrams of additional embodiments of the block diagrams of fig. 9.

Detailed description of some preferred embodiments of the invention

Making reference to fig.7, it is now described the statistical multiplexer scope of the present invention. Compared to the general description of fig. 2, the transmission 15 queues of QU cells grouping different connections, are here shared into m functional blocks addressed by the block indicators B-ID1, ..., B-IDm; indicators Q-ID are submitted to a similar sharing at the single transmission queues.

In fig. 7 we can notice two seemingly separate WFQ function, a first one for the processing of B-ID indicators, and a second one for the processing of Q-ID indicators. 20 A couple of indicators B-ID and Q-ID is generated any time a cell needs it. The B-ID indicator, according to the type of pre-set sharing between transmission queues of ATM cells and blocks, is, whenever absent, inserted in one of three FIFO lists, respectively identified in order of service priority with HP-FIFO, LP-FIFO, and VLP-FIFO. The indicator Q-ID, depending on the nature of the queue, which it is 25 associated to, is whenever absent, inserted in similar FIFO lists, associated to the block identified by the corresponding B-ID indicator. The two FIFO groups are "separately" used to perform the WFQ function in a way similar to that described, for the sole transmission queues, in the above mentioned patent application under the name of the same Applicant. The B-ID indicators selected by the first WFQ function are 30 therefore processed by a SHP/SKD block that performs the shaping and scheduling functions to limit the peak band of the set of transmission queues assigned to the generic block, afterwards the B-ID indicators time by time scheduled, are put at disposal of the second WFQ function for the selection of the group of three FIFOs to serve. Similarly to what done for the B-ID indicators, the Q-ID indicators supplied by the 35 second WFQ function are treated by an SHP/SKD block performing the shaping and

scheduling functions to limit the peak band of multiplexed flows.

Moreover, it must be highlighted that the VLP-FIFO lists are optional, as well as the SHP/SKD blocks.

In a first embodiment of the present invention it is foreseen an empty queue flag placed at the service of a queue grouping all the connections transmitting at the peak rate (CBR and rt-VBR) of a generic block addressed by B-ID, and requiring therefore to be served in real-time. The indication of said flag in the figures described up to now has been omitted, for sake of brevity.

In a second embodiment, all the connections transmitting at peak rate of all the blocks are grouped in a unique high priority block dedicated to the same, in this case both the WFQ functions are passed over and the Q-ID indicator is directly supplied to the relevant SHP/SKD block for its scheduling and the priority insertion of the cell on the output link. Also in this case, the indication of the empty queue flag associated to the unique queue grouping all the ATM cells of the real-time connections of all the blocks is missing in the figures elucidated up to now.

Making reference to fig. 8, it can be noticed that the generic block SHP/SKD, be it to the service of the transmission queues indicated by Q-ID, and by B-ID, includes a SHAPER-Q (SHAPER-B) block, performing a shaping function, placed upstream a SCHEDULER-Q (SCHEDULER-B) block, which in its turn performs a scheduling function. The shaping function calculates a waiting time, at the null limit, and converts the information in an N-SLOT-Q (N-SLOT-B) value that represents a slot of a calendar having finished length (indicated in the figure) included in the SCHEDULER-Q (SCHEDULER-B) block. The indicator Q-ID (B-ID) that reaches the shaper and the N-SLOT-Q (N-SLOT-B) value calculated by the shaper are transferred to the scheduling function that inserts the Q-ID (B-ID) indicator in the location of the calendar indicated by N-SLOT-Q (N-SLOT-B). The scheduling function performs the cyclic scanning of the calendar with $1/T_c$ interval, where T_c is the generic cell time, that is the time interval required to the serial transmission on the link of all the cell bits. At each location scanned, the indicators there stored are extracted from the calendar, on FIFO basis. In the figure, TNOW represents the actual cell time, that is the multiple time of T_c elapsed from the beginning of the scanning cycle. The time TNOW indicates in the figure an empty location, while the time TV indicates a previous calendar location.

The shaping function consists in limiting the queue and block band at top. For this reason it performs the calculation of the insertion time of Q-ID according to the following law:

INSERTION TIME = max[PREVIOUS(N_SLOT_QxTc) + 1/PCR, TNOW], where PREVIOUS(N_SLOT_Q) indicates the calendar location in which the Q-ID indicator was inserted the last time, and PCR (Peak Cell Rate) is the peak interval of the queue indicated by Q-ID. The same law applies to the blocks, provided that N-SLOT-Q is replaced by N-SLOT-B and a peak rate PCR is defined for each block. The configuration of PCR parameters in the SHAPER_Q and SHAPER_B blocks is made via software. As a consequence of the above mentioned insertion law, the scanning of the calendar will take care to appropriately delay the emission of queue and/or block indicators, arranging that the same are served not too quickly, thus violating the band assigned to the same.

Making reference to fig.9, it can be noticed how the system representation therein shown, differs from that of the known art of fig. 6 for the following:

- the presence of a WFQ-B block placed immediately upstream the SHP/SKD-B block, divided here into its components SHAPER-B and SCHEDULER-B;
- the WFQ-Q block is placed upstream the SHP/SKD-Q block, divided here into its components SHAPER-Q and SCHEDULER-Q;
- the indicator B-ID coming out from the SCHEDULER-B block is it too supplied to the buffer manager BM likewise Q-ID.

The functional block diagram of fig. 9 highlights the presence of two path rings of the indicators. The more internal ring is run across by the Q-ID indicators that start from the BM block and cross in order: WFQ-Q, SHAPER-Q, SCHEDULER-Q, and return to the starting BM block. The more external ring is run across by the B-ID indicators that start from the BM block and cross in order: WFQ-B, SHAPER-B, SCHEDULER-B, and return to the starting BM block. The second ring interacts with the first one at WFQ-Q block level. The reference to the rings has a mainly descriptive value.

Making reference to fig.10, it is now generally described the circuit architecture of the buffer manager BM of fig.9. As it can be noticed in the figure, the BM block includes the following sub-blocks: an access functionality controller ACC-CONTR, a controller of the service functionality SER-CONTR, a counter CODE-CONT associated to the transmission queues, a buffer for transmission queues BUF-CEL, a serial/parallel S/P converter, a parallel/serial converter P/S, and finally a distributor SYNC of the network clock CK. The controller blocks and the cell buffer communicate with a bi-directional data bus BUS-DATI, an address bus BUS-IND, and some control lines. The

input link LINK-IN reaches the converter S/P and from there, an input of the ACC-CONTR block, while the signal coming out from block SER-CONTR reaches the P/S converter to enter the output link LINK-OUT. The ACC-CONTR block sends a first writing enabling signal WR-EN to the BUF-CEL buffer for the writing of an ATM cell,
5 and a second signal WR-EN-P to a WFQ-CONTR unit for the execution of the WFQ algorithm (visible in fig. 11) for the writing of pointers B-ID and/or Q-ID associated to the relevant cell, having TR-TYPE value, characterizing the type of traffic supported by the connection which the cell belongs to, and a first and/or second count value VAL-CNT of two counters, respectively associated to the queue hosting the cell and to
10 the block hosting the queue. The SER-CONTR block sends a reading-enabling signal RD-EN to the BUF-CEL block. The counter CODE-CONT receives from the ACC-CONTR block a signal INCR-CONT to increase the count and a signal READ-CONT for the count reading, and a decrease count signal DECR-CONT from the SER-CONTR block. In response to the READ-CONT signal the counter gives the
15 datum read VAL-CNT. The SER-CONTR block receives at input two typologies of pointers (indicators) B-ID and Q-ID coming from the WFQ-Q and WFQ-B blocks (fig.9), or from SCHEDULER-Q and SCHEDULER-B blocks (fig. 9) in the case these last two blocks are not optional.

Before describing the operation of the buffer manager BM, partly already made
20 in the introduction, it is worth to describe figures 11 and 12 and then associate them to fig.10 in the operation.

Making reference to fig. 11, the circuit of blocks WFQ-Q and WFQ-B of fig. 9 is now generally described. As it can be noticed, and without any limitation for the invention, the architecture is referred to the case of control of the two blocks WFQ-Q
25 and WFQ-B through one sole execution unit of the WFQ algorithm, called WFQ-CONTR. The above mentioned unit is connected to the data bus and to the address bus of the buffer manager BM and receives from the ACC-CONTR block a writing enabling signal WR-EN-P. The unit -CONTR WFQ communicates with WFQ-Q and WFQ-B blocks through an internal bus BUS-DATI-WFQ and sends the
30 reading/writing signals to the above mentioned blocks, obtained from the decoding of the information received on the BUS-IND address bus or independently generated. The signals sent to the WFQ-Q block, are the following:

- W/R-ENH1, ...,W/R-ENHm; W/R-ENL1, ...,W/R-ENLm; W/R-ENV1, ...,W/R-ENVm;
- W/R-ENQ1, ...,W/R-ENQm;
- 35 • W/R-ENR1, ...,W/R-ENRm.

The signals sent to the block WFQ-B, are:

- W/R-ENH-B; W/R-ENL-B; W/R-ENV-B;
- W/R-ENB and W/R-ENR-B.

The WFQ-Q block includes m groups B_1, \dots, B_m , each one consisting of three memories FIFO, HP-FIFO, LP-FIFO, and VLP-FIFO. On the contrary, the WFQ-B block includes one sole group of memories FIFO. The WFQ-Q block includes also m RT-FLAG flags and m TK-CNT $_1, \dots, TK-CNT_m$ counters, whose function shall be clarified below, while the WFQ-B block includes one sole RT-FLAG flag and one single counter TK-CNT-B. Inside the WFQ-Q block, the reading/writing signals W/R-ENH reach the HP-FIFO memories, signals W/R-ENL reach the LP-FIFO memories, signals W/R-ENV reach the VLP-FIFO memories, W/R-ENR signals the RT-FLAG flags, and the W/R-ENQ signals the TK-CN counters in the corresponding numeral order. Inside the WFQ-B block reading/writing signals W/R-ENH-B, W/R-ENL-B, W/R-ENV-B reach the HP-FIFO, LP-FIFO, VLP-FIFO memories, respectively, the signal W/R-ENR-B the flag RT-FLAG, and the W/R-ENB signal the TK-CNT-B counter. The WFQ-CONTR unit placed on the internal data bus BUS-DATI-WFQ the pointers B-ID and Q-ID, values to write, or read, in the RT-FLAG flag and in TK-CNT counters selected at each cell time through development of the WFQ function.

Making reference to fig.12, it is shown the format of the j -th counter TK-CNT $_j$ among the counters TK-CNT $_1, \dots, TK-CNT_m$ of the WFQ-Q block, and of the counter TK-CNT-B of block WFQ-B. It must be said before that for the n queues inside a block the following relation applies: $z + s = n$, where z is the number of ABR connections, and s is the number of non-ABR connections. The counter TK-CNT $_j$ is a memory structure including z fields, N°TOK(QU $_1$), ..., N°TOK(QU $_z$), each one containing a number of "tokens", at disposal of each transmission queue QU $_1, \dots, QU_z$ belonging to the generic block B_j . Similarly, the counter TK-CNT-B includes m fields, N°TOK(B $_1$), ..., N°TOK(B $_m$), each one containing a number of tokens at disposal of a corresponding block B $_1, \dots, B_m$.

In the operation, what said about figure 2, reference to which shall be made, applies to the BM block, with the exception that B-ID block pointers add now to the sole Q-ID pointers. The previous illustration made in fig.2, considering the extremely general character of that figure, was not exhaustive concerning the destination of pointers. This gap is now filled, thanks also to what said concerning the WFQ function of fig.5. This said, the BM block performs the access function and ends the service one, essentially

developed by the WFQ-CONTR unit. The transfer of a pointer Q-ID from the ACC-CONTR control block to the WFQ-CONTR unit is necessary whenever at least one cell belonging to the considered queue is not already present in the BUF-CEL buffer. This information is contained in appropriate queue counters included in the count block CODE-CNT, each one of them indicating the number of cells belonging to a relevant queue present in the buffer BUF-CEL. The counters are updated by signals INCR-CNT and DECR-CNT, at each acceptance and service of the queue under examination, respectively. The transfer of B-ID is necessary whenever at least one cell belonging to a transmission queue assigned to the block considered is not already present in the BUF-CEL buffer. This information is obtained from block counters included in the CODE-CNT block, whose meaning as for the blocks is similar to that of the queue counters concerning the queues. The possible transfer of B-ID and/or Q-ID is completed with the traffic type TR-TYPE information. More in particular, there are three possible values of traffic type that originate a same number of destinations in the transfer of pointers Q-ID:

- HPT (High Priority Traffic): the pointer Q-ID is inserted in a relevant high priority memory HP-FIFO of the WFQ-Q, as to guarantee its privileged extraction reducing the staying time in the FIFOs of WFQ-Q blocks to the minimum extent. A queue characterized by a traffic type value equal to a HPT is suitable to contain connections CBR or rt-VBR.
- LPT (Low Priority Traffic): the Q-ID pointer is inserted in a relevant low priority memory LP-FIFO of the WFQ-Q until it has available tokens; once the consents are finished, the pointer is inserted in a very low priority memory VLP-FIFO. A queue characterized by a traffic type value equal to LPT is suitable to contain nrt-VBR connections of the ABR type (available band). This traffic type value, enables to assure a minimum band to the queues it characterizes.
- VLPT (Very Low Priority Traffic): the pointer Q-ID is inserted in a relevant memory VLP-FIFO at very low priority of the WFQ-Q. A VLP-FIFO memory is served only when the corresponding FIFOs having higher priority are empty. A queue characterized by a value of traffic type equal to VLPT is suitable to contain connections of UBR type, therefore a minimum band is not guaranteed.

On the contrary, concerning the sharing of B-ID pointers within the HP-FIFO, LP-FIFO, and VLP-FIFO memories of the WFQ-B block, the situation is not so well outlined as it is in the case of pointers Q-ID to the single queues. In fact, having to

decide in which FIFO a Q-ID pointer has to be stored, it is simply necessary to know the typology of the queue it is associated with, the information is supplied by the traffic-type parameter. The distinction between non-guaranteed minimum band flows and guaranteed minimum band flows, allows therefore assigning the latter to the 5 LP-FIFO whose service involves the management of the tokens. On the contrary, different queue typologies live together in a block, as for the service class. Therefore, it is necessary to define a criterion to share the B-ID pointers among the three FIFOs of the WFQ-B block. A first criterion could be that to subdivide the blocks (and therefore 10 their B-ID pointers) in three groups having equal aggregate band and to assign the groups to the FIFOs, according to the priority order, on the basis of the number of real-time queues contained in each group. This choice could cause an excessive increase of the CDV of the queues included in the blocks forming part of the group 15 associated to the VLP-FIFO memory, so that a second sharing criterion could no more consider the three groups as equiband, but as having increasing band in the sense of the decreasing priority. It is thus avoided the monopoly of the service by the group having higher priority, and, for the same reasons, by the group having intermediate priority versus the one having lower priority. Whichever is the block sharing principle, and leaving apart the second embodiment, it is convenient to arrange that the block 20 WFQ has the same operation of the queue WFQ, allowing an equal flexibility in the processing of blocks. Additional details shall be given when describing the operation flow-charts of the block WFQ.

The previous considerations on the conceptual and actual difficulty to subdivide the blocks into priority decreasing groups, together with the difficulty to share the transmission queues among the different blocks in order to make the selected priority effective, are clear indications of non obviousness of the present invention, to be opposed to possible conclusions of an analysis that, in a very reductive way, could see the present statistical multiplexing as the simple duplication of a known WFQ function. The absence of topics in favour of this wrong conclusion will also appear from the continued description of the operation and from the tenor of the claims, where it could 25 be noticed how the phases that make dynamic the statistical multiplexing of flows, that is the token assignment and renewal phases, are jointly performed for the blocks and for the relevant queues. This means that we are in presence of a more extended WFQ function capable of integrating the two procedures, complementary between them, reserved to the blocks and transmission queues, while maintaining separate the 30 phases that lead to the emission of queue block indicators. However, these phases are 35

not independent, since the phases leading to the emission of a transmission queue indicator can start only after a functional block has been selected, in this way the sharing of the residual band over the transmission queues results being conditioned. Too significant are the interrelations between the WFQ of queue block to think to a simple replication of a known algorithm; on the contrary it seems more logic to recognize the original character of the WFQ extended function used in the present invention.

Continuing in the description of the operation, we see that the acceptance of a cell from a transmission queue of cells extracted by the input flow LINK-IN (fig.1) terminates with the insertion of B-ID and/or Q-ID pointers in the appropriate FIFOs of the WFQ. The compliance with the insertion condition previously described assures that each pointer of queue blocks appears at most once in the FIFOs of the WFQ. The insertions of Q-ID and B-ID pointers, whenever necessary, are made in parallel.

The service of a transmission queue of cells to insert in the output flow LINK-OUT (fig.1) involves, at each cell time, the following operational phases developed in sequence:

- a) Extraction of a pointer B-ID from the group of FIFOs of the WFQ-B block (through execution of the WFQ algorithm);
- b) possible transfer of the B-ID pointer to the SHAPER-B block (fig. 9) and from this one to the SCHEDULER-B block to emit a B-ID pointer to a block with limited peak rate;
- c) extraction of a Q-ID pointer from one of the FIFO groups of the WFQ-Q block selected by the B-ID pointer emitted in the previous phase (through execution of the WFQ algorithm);
- d) possible transfer of the Q-ID pointer to the SHAPER-Q block (fig. 9) and from these to the SCHEDULER-Q block to emit a Q-ID pointer to a transmission queue having right to the service with limited peak rate;
- e) transfer of the B-ID and Q-ID pointers thus determined, to the buffer manager BM block to perform the following operations:
 - extraction of a cell from the BUF-CEL buffer by the SER-CONTR block for its insertion in the LINK-OUT output link;
 - generation of the DECR-CNT signal to decrease the block counters and queues of the counter block CODE-CNT;
 - generation of the READ-CNT signal for the reading of the VAL-CNT and of

CODE-CNT queue value of the block counters;

- possible generation of the WR-EN-P signal to re-write B-ID and/or Q-ID in the relevant FIFO group in case the relevant VAL-CNT value is not null.

As it shall be better seen in short, the B-ID and Q-ID pointers are emitted by
5 WFQ-B and WFQ-Q blocks according to the following law:

- if the HP-FIFO memory contains at least one pointer, it is extracted;
- if the HP-FIFO memory is empty and the memory LP-FIFO contains at least one pointer, it is extracted;
- if HP-FIFO and LP-FIFO memories are both empty and the VLP-FIFO memory
10 contains at least one pointer, it is extracted.

Making reference to figures 13, 14, 15, and 16, and to what said in the mentioned patent application filed by the same Applicant about the execution methods of the WFQ algorithm applied to the sole transmission queues, we now better specify the execution methods of the same in the frame of the statistical multiplexing method,
15 scope of the present invention, pointing out that in the new frame the above mentioned WFQ algorithm is applied in a novel way to select the blocks, and also to serve the relevant transmission queues.

The extension of the above mentioned algorithm also to queue transmission blocks is an absolute innovation of the present invention, and therefore the methods of
20 such an application shall be carefully described making reference to the main resources included in the WFQ-B block. With the appropriate distinctions, the assumptions described will continue to apply also to the WFQ-Q block.

The cases relevant to the first and second embodiments of the present invention, which foresee the use of the RT-FLAG to define an absolute priority on
25 real-time flows isolated in dedicated queues, or in a block, are a direct consequence of the complex of notions described, or mentioned. Therefore, a brief description shall be given only of the second embodiment, representing it too a novelty. Reference shall be made to what already said about the subdivision into service classes of the different flows entering a network node and multiplexed in the unique flow LINK-IN (fig. 1) at the
30 input of the generic statistical multiplexer SMU. It is worth to underline that there is not a prior relation between the band of the input link and the band of the output link LINK-OUT, and summarize that the input flow consists of cells belonging to a given number of flows with different QoS levels. These flows consist of a number z of

transmission flows of applications at available band or ABR (such as for instance the non "real time" flows relevant to a data file) for which the system assures in any case a minimum band and the maximum availability of the band in excess, and of a number s of transmission flows of non ABR applications (as are for instance the "real time" flows relevant to television signals, etc.), for which the system has however to guarantee the established peak band (CBR connections, that is maximum priority ones). To the above-mentioned flows, UBR type flows add, for which the service does not guarantee a minimum band. The mentioned QoS classes, together with what already said for the blocks, suggest the sharing criterion of B-ID and Q-ID pointers among the FIFO memories of the WFQ.

The assumption on which the WFQ algorithm bases, is that there is band available to share, therefore, denoted R_{tot} the total band available on the LINK-OUT link, denoted R_{Ai} the minimum bands negotiated of the z ABR transmission flows (or more in particular a minimum guaranteed band) and denoted R_{NAi} the fix transmission bands of the s non-ABR transmission flows (UBR excluded) the relation $\sum_i R_{Ai} + \sum_i R_{NAi} \leq R_{tot}$. has to apply. The observance of said relation is guaranteed, in the negotiation phase, by a correct assignment of the R_{Ai} and R_{NAi} values to the single applications.

The token mechanism was originated to assure the guarantee of a minimum band to ABR queues: To the purpose, the number of tokens assigned in a given TR period, called renewal period, to each transmission queue belonging to an ABR transmission flow, is proportional to the minimum band assigned in the negotiation phase. The renewal period TR is an integer multiple of the cell time Tc. It is defined as: TR = (R_{tot}/G) × Tc where the G factor characterizes the granularity tied to the minimum controllable band. The number of tokens assigned to each queue belonging to the ABR transmission flows is equal to R_{Ai}/G. The variation of the period TR, that is the variation of the granularity according to which tokens are distributed, enables to optimize the performances of the statistical multiplexer shown in fig.10 and 11 in any situation.

Profitably, the negotiated bands can easily be modified simply changing the number of tokens available for each application. In the negotiation phase, it is also possible to assign the tokens in such a way that the sum of all minimum bands of all ABR flows and of all peak bands of non ABR flows (UBR excluded) always results

lower than the total band of the output link. This will enable to always have available a given band in excess to share. In fact, the tokens can be seen as a band measurement unit, therefore the one who decides whether accepting or not a connection, shall verify that the traffic conditions in the network enable to assure to that connection and to the
5 other ones already present, the observance of the traffic parameters characterizing the same. If this evaluation is made using the token "measurement unit" then some tokens shall be assigned to each fixed band connection and to each connection with a guaranteed minimum band. From the above, it results that in the acceptance phase, some tokens are assigned also to the QUj queues whose pointers are contained in the
10 HP-FIFO memories of the WFQ-Q block. This same reasoning applies also to the blocks, since the band parameters of each block result from the bands of the relevant queues, therefore, in the acceptance phase, tokens are assigned also to the blocks whose B-ID pointers are contained in the HP-FIFO memory of the WFQ-B block. Being valid the criterion for the assignment of the tokens established for the queues, it results
15 that the sum of the tokens of all the blocks does never exceed the total output band (LINK-OUT).

However, for the statistical multiplexer according to the present invention it remains valid the fact to use the tokens only to assure the minimum bands, and for this reason the fixed band connections are always served with high priority, irrespective of
20 the value of the tokens assigned to each one of them by the one who decides the acceptance. Now the question arises spontaneous: and if a fixed band connection (e.g. CBR) would transmit at a band higher than its PCR (Peak Cell Rate), should it not actually receive more than due? The answer is affirmative, but this cannot happen since, as it is known, some components called "policers" (not indicated in the figures)
25 are present in the input nodes of the network (fig.1) that check and arrange that each connection CBR (VBR) does not transmit at a band higher than its PCR (SCR).

Making reference to fig.13 it is described the method used in the present invention for the sharing of the available band over the different B-ID blocks. The method is described concerning one sole renewal TR period.

30 At the beginning of the TR period the control unit WFQ-CONTR (fig. 11) sets on an INIZIALIZZA TR-CNT phase in which a TR-CNT counter is initialized with the TR value of the token renewal period.

An INIZIALIZZA TK-CNT-B, TK-CNT1,...,TK-CNTm phase follows, during which the storing devices of the above mentioned tokens (fig.12) are initialized with token
35 values N°TOK(B1), ...,N°TOK(Bm) assigned to the relevant fields of the TK-CNT-B

counter, and with token values $N^{\circ}TOK(QU_1), \dots, N^{\circ}TOK(QU_z)$ assigned to the relevant fields of the m counters $TK-CNT_j$. The number of tokens $N^{\circ}TOK(B_j)$ that a generic block actually employs to assure the minimum block band is given by the following expression:

$$5 \quad N^{\circ}TOK(B_j) = \sum_{i=1}^z N^{\circ}TOK(QU_i)$$

where z is the number of transmission queues QU of ABR type assigned to the j -th block.

Thereafter, the WFQ-CONTR unit sets on a INITIALIZE FIFO OF WFQ-Q, WFQ-B BLOCKS phase, during which possible pointers contained in the above mentioned memories are removed, and opportunely inserted in the groups of HP-FIFO, LP-FIFO, VLP-FIFO memories of the WFQ-Q block of transmission queues Q-ID pointers having cells to transmit, and in the memories, bearing the same name, of the WFQ-B block of B-ID pointers to the blocks including the above mentioned queues.

The initialisation phases described up to now, clearly show how up to this point 15 two independent WFQ functions, for the blocks and the transmission queues, respectively, cannot be identified at all, since, in fact, the relevant processing is jointly made. From this point on, it is possible to separately describe the two functions, but this does not mean that the second one is separate from the first one.

The WFQ-CONTR unit in the TR-CNT – $T_c \rightarrow TR-CNT$ phase decreases by a 20 quantity corresponding to the cell time T_c the counter TR-CNT, and sets on a test phase $TR-CNT = 0$ in which it checks if the time TR is not out.

If the time TR has not yet elapsed, the WFQ-CONTR unit sets on a reading phase READ HP-FIFO during which a B-ID pointer is read, followed by a test phase IS 25 HP-FIFO EMPTY? in which it checks if in the above-mentioned FIFO there are no more service high priority block B-ID pointers still to emit.

If the test indicates that said pointers are present, the unit WFQ-CONTR sets on a phase SERVI HP-FIFO during which the block whose B-ID pointer had been read in the READ HP-FIFO phase, is served. On completion of the SERVI HP-FIFO phase, that shall be described in association with fig. 14, the WFQ-CONTR unit returns to the 30 phase in which TR-CNT is decreased and cyclically repeats the phases included between this last and the service phase until the time TR is not out, or until the HP-FIFO memory itself has not been emptied.

When this occurs, the test IS HP-FIFO EMPTY? results true and the unit WFQ-CONTR sets on a reading phase READ LP-FIFO during which a pointer B-ID is

read, followed by a test phase IS LP-FIFO EMPTY? In which it checks if no more service low priority B-ID pointers still to emit are present in the above mentioned FIFO.

If the test indicates that said pointers are present, the WFQ-CONTR unit sets on a phase SERVI LP-FIFO during which it is served the block whose B-ID pointer has been read in the READ LP-FIFO phase. At the end of the SERVI LP-FIFO phase, which shall be described in conjunction with fig. 15, the WFQ-CONTR unit returns to the phase in which TR-CNT is decreased and cyclically repeats the reading and service of the LP-FIFO memory phases, until the TR time is not out, or until the LP-FIFO memory itself has not been emptied.

When this occurs, the test IS LP-FIFO EMPTY? results true and the WFQ-CONTR unit sets on a reading phase READ VLP-FIFO during which a pointer B-ID is read, to which a test phase IS LP-FIFO EMPTY? follows (transferred for representation convenience at the beginning of fig. 16) in which it checks if no more very low priority service blocks B-ID pointers still to emit are present in the above mentioned FIFO.

If the test indicates that said pointers are present, the WFQ-CONTR unit sets on a phase SERVI VLP-FIFO during which the block whose B-ID pointer was read in the READ VLP-FIFO phase is served. At the end of the SERVI VLP-FIFO phase, that shall be described with fig. 16, the WFQ-CONTR unit returns to the phase in which TR-CNT decreases and cyclically repeats the reading and service phases of the memory VLP-FIFO until the time TR has not elapsed, or until the VLP-FIFO memory itself has not been emptied. In both the cases the START phase is restarted and the WFQ-CONTR unit newly sets on the initialization phases of the counter TR-CN in which the original value of the renewal period TR in the counter TR-CNT is restored, followed by the phase in which the tokens in the counters TK-CNT₁, ..., TK-CNT_m and TK-CNT-B, FIFO are newly calculated and redistributed, and the phase in which the content of all the FIFOs of WFQ-Q and WFQ-B blocks are re-initialized.

Making reference to fig. 14, it is described the SERVI HP-FIFO phase of fig. 13 which starts with an EMETTI B-ID SU BUS-DATI-WFQ phase whose meaning is obvious.

A phase follows in which the VAL-CNT value is read, referred to the needs of the block indicated by B-ID, and a test is made on the value read to define whether blocks of that kind still to be served are present. This shall be true if at least one cell still to transmit is present in at least one queue assigned to the block just served.

If the test VAL-CNT = 0? results false, the phase WRITE B-ID IN HP-FIFO is

run, to reinsert the B-ID pointer in the HP-FIFO memory, otherwise the WFQ-CONTR unit sets again on the time decrease TR phase and repeats the emptying cycle of HP-FIFO, thus terminating the service phase.

Making reference to fig. 15, it is now described the SERVI LP-FIFO phase of fig. 13 that starts with a phase EMETTI B-ID SU BUS-DATI-WFQ whose meaning is obvious. Then the VAL-CNT value is read, referred to the needs of the block indicated by B-ID, and a test is made on the value read to define whether blocks of that kind still to be served are present.

If the test $VAL-CNT = 0$ results false, the WFQ-CONTR unit performs in the phase IS TK-CNT-B = 0 a test on the value of the token counter TK-CNT-B field, referred to the block indicated by the pointer B-ID.

If from the test it results that the subject block still has tokens to spend, the pointer B-ID is re-entered in the LP-FIFO memory and the value in the field of the token counter TK-CNT-B is decreased by one unit. After that, the WFQ-CONTR unit sets on the time decrease TR phase and repeats the emptying cycle of LP-FIFO.

If from the test on the token counter TK-CNT-B it results that the block indicated by B-ID has already spent all its tokens, the pointer B-ID is inserted in the VLP-FIFO memory, afterwards it is decreased by a time unit TR and the emptying cycle of LP-FIFO repeated, thus terminating the service phase.

Fig. 16 shows the phase SERVE VLP-FIFO of fig.13 that is completely similar to the SERVE HP-FIFO phase of fig.14, except for the fact that the pointer B-ID is now entered in the VLP-FIFO memory.

It is absolutely clear that, temporarily leaving apart the subsequent scheduling operations of B-ID and Q-ID indicators, once an EMETTI B-ID SU BUS-DATI-WFQ phase results completed, it can be started the execution of those phases of the WFQ applied to the transmission queues, leading to the emission of a Q-ID indicator and which are completely similar to the corresponding phases described for the blocks.

From the flow charts just described, we can summarize the main characteristics of the WFQ block, stating that the high priority FIFO, HP-FIFO contains functional block B-ID pointers that are always served, irrespective of the tokens assigned to the same during the negotiation phase.

The low priority FIFO, LP-FIFO, contains pointers B-ID at functional blocks still having token available and that, therefore, have certainly right to be served.

The very low priority FIFO, VLP-FIFO, contains functional block B-ID pointers that have already spent all the tokens at their disposal and that have therefore certainly

fully used the minimum band allowed to the same. Therefore, said blocks have equal right to be served, so the residual band is evenly shared among the same.

During the time interval TR the memory LP-FIFO is cyclically scanned and the blocks indicated by the same are evenly served. These blocks, due to the decrease of 5 the number of tokens and of their call for the service of the relevant transmission queues, can be in different states.

A block can simultaneously finish its token and have no queue to serve left. This situation shall occur when, during the TR time, the aggregate cell-rate has been just equal to the minimum band assigned to the block.

10 On the contrary, a block can also have no more queues to serve before having finished its tokens. This means that this block will include transmission queues whose aggregate cell-rate, during the TR period, has been lower than the minimum band assigned to the block.

Finally, one block can have no token left before the timeout of the TR period. 15 This means that this block shall include transmission queues whose aggregate cell-rate, during the TR period, has been higher than the minimum band assigned to the block.

During the TR period, due to the emptying activity of the LP-FIFO memory, the following two cases could occur:

20 In the first case, the emptying of the LP-FIFO coincides with the end of the TR period. This means that all the functional blocks have transmitted, during the TR period, with an aggregate cell-rate equal to the speed determined by the minimum band agreed. In this case, during such period, there will be no band available in excess to distribute among the blocks.

25 In the second case, the emptying of the LP-FIFO occurs before the end of the TR period. This means that one at least among the functional blocks has transmitted, during the TR period, with an aggregate cell-rate lower than the speed determined by the agreed minimum band. In this case, during this period, there will be available band in excess to share among the blocks. The VLP-FIFO memory is cyclically scanned in 30 such a way that all the blocks whose indicator is present are evenly served, irrespective of the number of tokens owned in origin by them and irrespective of the number of queues they have still to serve (Round-Robin).

It could be appreciated that, in this way, even in the worst case in which all the 35 blocks have queues whose aggregate cell-rate is higher than the minimum band assigned to the relevant block, the blocks themselves are correctly served at least with

the minimum band agreed. It can be also appreciated that, thanks to the scanning method of the VLP-FIFO memory, no block can monopolize the service.

In the specific case of the second embodiment, in addition to the three FIFO memories, the flag RT-FLAG is used to govern the unique block grouping all the transmission queues with real-time service needs. In this case the WF-CONTR unit, once the initial phase is completed, immediately proceeds to perform a test on the RT-FLAG value to verify if there is a block having absolute priority to serve, even compared to the high priority FIFO, HP-FIFO. Once the above mentioned block is emptied and the flag is disabled, the WF-CONTR unit proceeds as shown in fig. 13.

Making reference to fig.8, it is now discussed the case in which the SHAPER-B, SCHEDULER-B and SHAPER-Q, SCHEDULER-Q blocks of fig. 9 have not to be considered as optional, but effective. This is the most likely case since, generally the blocks and the transmission queues with traffic-type non real-time require a prior control of the peak band before they are admitted to the service. We remember that for the corresponding real-time the peak is guaranteed in advance in the connection acceptance phase. In the case considered, the main flow chart of fig. 13, as well as the flow charts of the service phases shown in figures 14, 15 and 16 remain unchanged, since the service they render is to emit a pointer B-ID and/or Q-ID on the data bus BUS-DATI-WFQ of the WFQ-CONTR control unit. The above mentioned unit shall decide, from the analysis of the TR-TYPE, whether extending the bus BUS-DATI-WFQ towards the SHAPER-B and SHAPER-Q blocks, or transfer the pointer B-ID and/or Q-ID directly to the SER-CONTR unit (fig. 10) controlling the terminal phase of the service functionality. In a completely similar way, it is possible to selectively disable, via software, the peak control on the band of single blocks, and/or on the band of single transmission queues, which shall continue to have a band controlled by the pure WFQ algorithm. The proposed structure enables therefore the coexistence of blocks and/or queues with or without control of the peak cell-rate.

Fig. 17 shows the functional block diagram of a third embodiment in which the search for a B-ID indicator to emit is made in parallel by SCHEDULER-B blocks and by the tandem WFQ-B, SHAPER-B, and the search for a Q-ID indicator to emit is made in parallel by the SCHEDULER-Q blocks and by the tandem WFQ-Q, SHAPER-Q. The resulting advantage is to reduce the CDV suffered by the indicator when completely running along its ring.

Making reference to fig.17, it can be noticed how the block diagram shown therein, differs from that of fig. 9 in the following:

- the output of the SHAPER-Q block is no more connected to the SCHEDULER-Q block but to a first input of a SELECT Q-ID block;
 - the output of the SCHEDULER-Q block is no more connected to the buffer manager BM but to a second input of the SELECT Q-ID block;
- 5 - a first output of the SELECT Q-ID block is connected to the buffer manager BM;
- a second output of the block SELECT Q-ID is connected to the input of the SCHEDULER-Q block.

Dual considerations apply to the WFQ-B, SHAPER-B, SCHEDULER-B blocks; in this case a SELECT B-ID block is present, whose output connected to the buffer manager BM is also connected to a third input of the WFQ-Q block. For the sake of brevity, we have neglected in the figure the indication of a relevant control unit of blocks SHAPER-Q, SCHEDULER-Q, and SHAPER-B, SCHEDULER-B.

Making reference to fig. 18, it is now described the operation of the statistical multiplexer represented by the block diagram of fig. 17 concerning the sole shaping and scheduling functions. For the sake of brevity, it is indicated the flow chart of the above mentioned functions concerning B-ID indicators only, being the flow chart concerning the indicators Q-ID completely similar to the one that shall be now described. In fig. 18 after the START indication, the processing flow divides into two branches made in parallel. On the first branch, the B-ID indicator extracted from the WFQ-B block is sent, at the end of the ESEGUI WFQ-B phase, to the SHAPER-B block. This last calculates the insertion time in the calendar, expressed as N-SLOT-B, and in the subsequent phase the control unit inserts B-ID in the calendar at the location indicated by N-SLOT-B, increases by a cell time Tc an actual time counter TNOW used for the scanning of the calendar, and returns to the starting point, immediately after the START. Observing the block diagram of fig. 17, the B-ID → CALEND(N-SLOT-B) phase is equivalent to the path made by B-ID by the second output of the SELECT B-ID block at the input of the SCHEDULER-B block.

Simultaneously to the above-mentioned processing phases, the control unit performs a test phase È B-ID DA SCHED-B to check the presence of an indicator B-ID emitted by the SCHEDULER-B.

If the response is yes, the B-ID pointer is emitted on the first output of the SELECT B-ID block from where it proceeds towards the buffer manager BM, and the control returns to the execution of the TNOW+Tc → TNOW phase and to the subsequent repetition of the processing on the two branches.

35 If on the contrary the response is no, it means that the research algorithm

governing SCHEDULER-B has not identified any B-ID indicator to extract, the control unit performs therefore a second test phase È B-ID DA WFQ-B to check the presence of an indicator B-ID emitted by the WFQ-B block.

5 If the response is no, no B-ID indicator has been emitted by WFQ-B. In this case no B-ID indicator (and therefore Q-ID) can be emitted by the ring in the present cell time and no additional operation shall be made, therefore the control returns to the execution of the TNOW+Tc → TNOW phase and to the subsequent repetition of the processing on the two branches.

10 If in the second test phase the response is yes, it means that a B-ID indicator has been emitted by the WFQ-B block and it must be defined if and when it could be re-emitted towards the BM block. To this purpose, a third test phase È N-SLOT-B > TNOW/Tc is performed, in which the control unit of the scheduler verifies if the insertion location in the calendar calculated by the SHAPER-B block is higher than that (TNOW/Tc) corresponding to the present time TNOW.

15 In the affirmative, no B-ID indicator can be extracted from the ring in the present cell time TNOW, and the control proceeds to a subsequent phase INSERT B-ID IN CALENDAR AT N-SLOT-B LOCATION, whose meaning is obvious; thereafter the control returns to the execution of the phase TNOW+Tc → TNOW and to the subsequent repetition of the processing on the two branches.

20 If in the third phase the response is no, it means that the location N-SLOT-B is equal to that calculated by the shaper at the actual time TNOW, the sign of minor would represent in fact a condition made impossible by the particular algorithm adopted by the shaper. In this case, it is performed a phase EMETTI B-ID in which the indicator B-ID is emitted on the first output of the block SELECT B-ID, from where it proceeds towards the buffer manager BM. No insertion of B-ID in the calendar is required. Thereafter, the control returns to the execution of the TNOW+Tc → TNOW phase and to the subsequent repetition of the processing on the two branches.

30 Profitably, one single integrated circuit of the ASIC type (Application Specific Integrated Circuit) has been implemented, including all the functional blocks indicated in fig.9.

Architectures with reduced functionality can be obtained removing from the architecture of fig. 9 some blocks that implement the functions not considered. In this case, it shall be necessary to consistently modify the flow charts described up to now. In the multiplexer shown in the third embodiment of fig. 17 for instance, it is possible to

include the parallel processing offered by the SELECT block in one or in the other, or in both, the rings of emission of pointers B-ID and Q-ID.

Another example of these structures can be seen in fig. 19 concerning a fourth embodiment of the present invention, where SHAPER-Q and SCHEDULER-Q have
5 been removed. In this case the queue indicator Q-ID extracted from the WFQ-Q block is not inserted in the calendar but directly sent to the buffer manager BM, which shall request the reading of an ATM cell stored in the buffer BUF-CEL (fig. 12) and, when necessary, re-transfer the indicator Q-ID to the WFQ function-Q. The solution indicated in fig. 19, that does not allow checking the peak band of each transmission queue,
10 however assures a better use of the band and a lower occupation of logic.

A further example of structure with reduced functionality is visible in fig.20, concerning a fifth embodiment of the present invention in which the WFQ-B block has been removed. In this case the BM block transfers the B-ID indicators directly to the SHAPER-B block which, on completion of its function, inserts it in the calendar scanned by the SCHEDULER-B block. This choice is functionally equivalent to the operation method of the known art, since it assigns a constant band to the blocks and reserves the application of the WFQ function to the sole transmission queues.
15